

AD-A192 665


HUMAN ORGANIZATIONS AS DISTRIBUTED INTELLIGENCE SYSTEMS 1/1
(U) MASSACHUSETTS INST OF TECH CAMBRIDGE LAB FOR
INFORMATION AND DECISION SYSTEMS A M LEVIS FEB 88


UNCLASSIFIED

LIDS-P-1752 N00014-84-K-0519

F/G 12/9

NL





MICROCOPY RESOLUTION TEST CHART
URBAN U. STANDARDS-1963-A

AD-A192 665

February 1988

④
ONE FILE COPY

LIDS - P - 1752

HUMAN ORGANIZATIONS AS DISTRIBUTED INTELLIGENCE SYSTEMS*

by

Alexander H. Levis **

ABSTRACT

Human decision making organizations can serve as a paradigm for distributed intelligence systems. Elements of a mathematical theory of organizations are presented and used to analyze the behavior of organizations that have to meet stringent requirements with constraints on the cognitive and physical resources available. It is shown how the introduction of decision support systems to aid individual decisionmakers can affect in unforeseen ways system performance.

DTIC
SELECTED
S MAR 24 1988
D
S
H

* This work was carried out at the MIT Laboratory for Information and Decision Systems with support provided by the Office of Naval Research under contract No. N00014-84-K-0519 (NR 649-003).

** The author is with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.

DISTRIBUTION STATEMENT A
Approved for public release
Distribution Unlimited

88 3 10 044

INTRODUCTION

Human organizations are by definition distributed intelligence systems. The characterization applies regardless of which definition is used for the terms distributed and intelligence. A common dictionary defines intelligence as the capacity for reasoning, understanding, or for similar forms of mental activity - clearly, human characteristics. Distributed means that some resource - intelligence in this case - is dispersed through a space or an area.

A more interesting set of definitions is given by Minsky (1986). He defines distributed processes as those in which each function is spread out over a range of locations, so that each part's activity contributes a little to each of several different functions. Substitute "organization member" for "part" and the relevance of the definition is clear. His definition of intelligence, however whimsical, is apt when applied to human organizations - it is all the mental skills that, at any particular moment, we admire but don't yet understand.

In this paper, a mathematical theory of distributed decisionmaking organizations will be presented with emphasis on the concepts, on some of the insights obtained, and on many of the challenges that remain. Drenick (1986) states in his recent book on a mathematical organization theory that the objective of such a theory is "to derive certain conclusions from a set of assumptions by mathematical reasoning." There are two consequences of this statement.

The first is that the assumptions must be stated explicitly and unambiguously. In the beginning, the assumptions are, by necessity, very restrictive since one is dealing with such a complex phenomenon as distributed intelligence in an organization. Then, one carefully and systematically proceeds to weaken the restrictive assumptions so that more aspects of organizational performance can be captured.

The second consequence is that this approach can lead to model-based experimentation: the mathematical theory can give rise to a set of testable hypotheses around which controlled experiments can be designed. With this empirical evidence and the mathematical theory, one can then begin to address the design problem. The nature of the problem is such that a synthesis is required of concepts, methods, and tools from such fields or disciplines as system theory, information theory, computer science, management science, and cognitive and behavioral psychology.

See Letter

Distribution/	
Availability Codes	
DIST	Avail and/or Special
R-1	

Galbraith (1977) describes the various approaches to organization design as follows: In classical management theory, the key concept is that of division of labor. A modern counterpart of that approach is the functional decomposition of tasks and the subsequent allocation of those to the organization members so that some objective function is maximized. Parallel, hierarchical, and mixed structures result. In contrast to this mechanistic approach is the human relations approach in which the style of leadership determines the organizational form. Empirically based, it leads to the consideration of incentives - tangible and intangible - as the means to improve performance. The third approach, which is closest to the conceptual framework for distributed intelligence systems, is based on the view of the organization as an information processing and decision making system. The cognitive limitations of humans (or the memory and information processing limitations of machines) determine the organizational form. This approach admits the allocation of functions to humans and machines. If the decision problems can be formulated analytically and if the proper data are available in the proper form, then the solution can be obtained algorithmically. In the idealized case, classical rationality is used.

But, as more operational problems under more realistic conditions are considered, the classical rationality model does not hold. Even if a clearly defined, commonly agreed upon objective function is identified, there is never enough time to enumerate the possible options, evaluate them, and select the best. Time constraints are usually the most potent reason for violating classical rationality. More subtle reasons also arise. For example, the data available for the evaluation of two competing options may not be consistent and concurrent, thus making comparison not strictly valid.

As a result, to improve organizational performance, decision aids have been introduced that sometimes aim at reducing the decisionmaker's workload by carrying out mundane, but time consuming tasks such as evaluation of alternatives; sometimes aim at augmenting the decisionmaker's scope by introducing additional options; and sometimes aim at reducing human error by providing guidance through a checklist or a step by step procedure. In all cases, they have complicated the organization design problem significantly. The information processing and decision making functions are now distributed not only among the intelligent organization members, but also between humans and machines. This is the design challenge posed by considering human organizations as distributed intelligence systems.

Key assumptions are presented first followed by the model of the organization member. In the

next section, organizations will be constructed. Then measures of performance will be described and the models used in their evaluation outlined. Finally, some of the results obtained to date will be discussed - in particular, consequences of the distributed nature of the cognitive processes.

ASSUMPTIONS

A restricted class of organizations will be considered. It is assumed first that the organization consists of at least two human decisionmakers and that it is a team. A team is defined as an organization in which the members have a common goal, have the same interests and same beliefs, and have activities that must be coordinated so as to achieve a higher effectiveness (Grevet, 1987). It is further assumed that they are well trained for the tasks that they have to perform and that they do not learn during the execution of a particular task.

It should be possible to draw a boundary that defines what is included in the organization and what is excluded, i.e., what resides in the external environment. Tasks that the organization must perform are generated in the environment by one or more sources which may or may not be synchronized. The organization acts upon these inputs and produces a response, including the null response, that is directed to the environment. Thus, the interface between the system and the environment is composed of the sensors and the effectors. Whether to include the sensors or the effectors or both as parts of the organization or as parts of the environment is a question that must be addressed in each particular case. This issue becomes relevant when alternative organizational designs are evaluated; the comparison must be done on the same basis.

The elements of the organization consist of the human decisionmakers, data bases, processors, and communication systems. A decision aid is defined as any technique or procedure that restructures the methods by which problems are analyzed, alternatives developed, and decisions taken. Decision support systems, a specific form of decision aids, do not automate a specific decision making process, but must facilitate it (Keene and Scott Morton, 1978). Decision support systems are considered here as higher level components that may consist of processors, data bases and communication systems.

Relationships are the links that tie these elements together. These relationships can be considered at three levels: they may describe the physical arrangement of the components - such as the

geographical location of the organization members, - or the functional relationship between components - such as the sharing of information between two members, - or the rules and protocols that govern the interactions - such as the conditions under which two members may share information. While this demarcation between relationships and components is often hard to justify, it is assumed that it can be done.

THE ORGANIZATION MEMBER

The model of the human decisionmaker who interacts with other decisionmakers and with the environment is shown in Fig. 1 (Levis, 1984).




Fig.1. The interacting desisionmaker with memory.

The decisionmaker (DM) receives input signals x from a variety of sources: from the environment, from a decision support system (DSS), or from the rest of the organization. He can receive one input at a time at the Situation Assessment (SA) stage. He processes this input, with or without use of information stored in a data base (memory) to obtain an estimate of x , the "assessed situation" z , which he may share with other DMs. He may also receive at this point other information, z'' , from the rest of the organization.

He combines this information with his own assessment in the Information Fusion (IF) stage, which contains a data fusion algorithm, to obtain the final assessment of the situation, labeled z' . The next step is the consideration of commands v' from other DMs which could result in a

restriction of his set of alternatives for generating the response to the given input. This is the Command Interpretation stage, or CI. The outcome of the CI stage is a signal v which contains the data z' and the rule v' used in the Response Selection (RS) stage to select the procedure or algorithm for generating the output y . This is the response of the decisionmaker; it may be sent to the external effectors or to other DMs within the organization.

The Petri Net formalism (Peterson, 1981) has been found very convenient for describing the concurrent and asynchronous characteristics of the various interactions. Petri Nets are bipartite directed multigraphs. The two types of nodes are the places, which represent signals or conditions, and the transitions, which represent processes or events. Places are denoted by circles and transitions by bars. A marking of a Petri Net assigns a non-negative integer number of tokens to each place. A transition is enabled, if and only if each of its input places contains at least one token. The places act like buffers, hosting the tokens until all the input places of a transition are non-empty. Enabled transitions can fire. When they fire, a token is removed from each input place, and a token is deposited in each output place. In the Petri Net representation of the DM model, the transitions stand for the algorithms, the connectors for the precedence relations between these algorithms, and the tokens for their input and output. The time taken by the algorithm to run is the transition processing time $\mu(t)$. The Petri Net model of the four stage DM without memory is shown in Fig. 2.




Fig. 2. Four stage model of a DM.

The tokens, in the simplest version of the model, are all indistinguishable. A token in a place means simply that an item of information is available to the output transition(s) of that place. It is also possible to associate attributes with the tokens. In this case, the source can be represented

by a finite number of distinct tokens x , each one occurring with some probability $p(x)$. However, if the protocols ruling their processing do not vary from one set of attributes to the other, they can be considered as indistinguishable tokens.

The intelligence in this model is embodied in the algorithms imbedded in the transitions; however, even if the algorithms are stochastic, the model is rather mechanistic and does not capture human decisionmaking well. To model the choice inherent in decisionmaking, a decisionmaker is assumed to have, at any stage of his processing, a set of options: different algorithms processing in different ways the same input to produce the same type of output. Thus, the SA and RS stages of the decisionmaker of Fig. 2 are modeled so as to include a set of U and V algorithms, respectively. The SA and RS stages are represented by a subnet of a Petri Net with switches (Fig. 3).




Fig. 3. Four stage model with switches.

Switches are transitions which resolve conflict situations; a switch is a transition with multiple output places and a rule according to which one and only one of the output places is chosen to receive a token after the transition has fired. In the SA stage, this choice is denoted by the variable u , taking its values in $\{1, 2, \dots, U\}$. The rule for determining the value of the decision variable is called the decision strategy of the decisionmaker for the particular stage. If the rule is characterized by a probability distribution $p(u)$ and if one branch of the switch is always chosen, i.e., if there is an i in $\{1, \dots, U\}$ such that $p(u = i) = 1$, then the strategy is called pure. Otherwise, it is mixed. The strategy that a decisionmaker uses at the RS stage usually depends on

the input to that stage. In that case, the probabilities are conditional probabilities $p(v=j | z, v)$. Together, the strategies for the two stages constitute the internal decision strategy of the DM. While this is a way of describing the set of strategies that a well trained decisionmaker may use, if his bounded rationality threshold is not exceeded, there are no rules to specify how and when any of these strategies will be selected by a specific decisionmaker at any given time. These rules are assumed to depend on the level of expertise of the DM, and on those mental skills that "we admire but don't yet understand," i.e., on the DM's intelligence.

The workload of each decisionmaker reflects the mental effort required to carry out the information processing and the decision making. A mathematical model of the workload has been developed (Boettcher and Levis, 1982) that is based on n - dimensional information theory. It is based on the assumption that the higher the uncertainty in the input, the more processing has to be done to reduce uncertainty to the point that a decision can be made. The value of the workload G is obtained by computing the entropy of all the internal variables of the DM model. The cognitive limitations of human DMs can be modeled in terms of the bounded rationality constraint. This is based on the premise that the rate with which decisionmakers process information is bounded; if the rate is exceeded, then rapid degradation of performance occurs. Formally,

$$G / \tau \leq F_{\max}$$

where F_{\max} is the maximum rate and τ is the mean interarrival time. A recent experiment at MIT (Louvet et al., 1987) has shown that for well defined cognitive tasks, F_{\max} exists, is stable, and is normally distributed across decisionmakers.

With this model of the organization member, it is now possible to formulate the problem of designing distributed decisionmaking organizations.

ORGANIZATIONS

Interactions between DMs

It was shown in Figs. 2 and 3 that a decisionmaker can only receive inputs at the SA, IF, and CI stages, and produce outputs at the SA and RS stages (Remy et al., 1988). These conditions lead

to the set of admissible interactions between two DMs that is shown in Fig. 4. For clarity, only the connectors from DM^i to DM^j are shown; the interactions from DM^j to DM^i are identical.




Fig. 4. Allowable interactions between DMs.

The mathematical representation of the interactions between DMs is based on the connector labels e_i , s_i , F_{ij} , G_{ij} , H_{ij} , and C_{ij} of Fig. 4; they are integer variables taking values in $\{0,1\}$ where 1 indicates that the corresponding directed link is actually present in the organization, while 0 reflects the absence of the link. These variables can be aggregated into two vectors e and s , and four matrices F , G , H , and C . The interaction structure of an n -decisionmaker organization may be represented by the following six arrays:

Two $n \times 1$ vectors e and s , representing the interactions between the external environment and the organization:

$$e \equiv [e_i]; \quad s \equiv [s_i]; \quad \text{for } i = 1, 2, \dots, n.$$

Four $n \times n$ matrices F , G , H , C representing the interactions between decisionmakers inside the organization:

$$F \equiv [F_{ij}]; \quad G \equiv [G_{ij}]; \quad H \equiv [H_{ij}]; \quad C \equiv [C_{ij}]$$

for $i = 1, 2, \dots, n$ and $j = 1, 2, \dots, n$.

Since there are four possible links between any DM and any other DM except himself, then the

maximum number of interconnecting links that an n -decisionmaker organization can have is

$$k_{\max} = 4n^2 - 2n.$$

Consequently, if no other considerations were taken into account, there could be

$$2^{k_{\max}}$$

alternative organizational forms. This is a very large number: 2^{90} for a five person organization. Fortunately, it can be shown that the set of all nets described by the six arrays is a partially ordered set and, furthermore, that it is a lattice (Remy and Levis, 1988).

Generation of Architectures

The analytical description of the possible interactions between organization members forms the basis of an algorithm that generates all the architectures that meet some structural constraints as well as application-specific constraints that may be present. The set of structural constraints that has been introduced rules out a large number of architectures. The most important constraint require that a directed path exist from the source to every node and from every node to the sink, and that the organizational structure be acyclical. The first constraint addresses the connectivity of the organization - it eliminates structures that do not represent a single integrated organization.

An algorithm has been developed (Remy and Levis, 1988) that determines the maximal and minimal elements of the set of designs that satisfy all the constraints; the entire set can then be generated from its boundaries. The algorithm is based on the notion of a simple path - a directed path without loops from the source to the sink. Feasible architectures are obtained as unions of simple paths.

In more recent work by Andreadakis and Levis (1987) a slightly different model has been used. Instead of using the four-stage model of the decisionmaker, a five-stage information and decisionmaking process has been defined. This constitutes the data flow structure that can be partitioned by allocating the stages to different decisionmakers. This approach has allowed the introduction of two additional design specifications: the degree of complexity of the organization and the degree of redundancy. The first measure addresses the complexity that results from

transitions needing many different inputs to be enabled. The second reflects the number of output places associated with each transition.

Thus, it has become possible to generate a complete class of distributed architectures that can serve as candidate structures. The reduction of the computational complexity of the problem to a tractable level is due to the presence of constraints and on the special structure of the decisionmaker model.

Variable Structure Organizations

The set of interactions between decisionmakers in the model considered thus far does not depend on the position of the switches imbedded in the situation assessment and response selection transitions (see Fig. 3). The interactions are always there and, consequently, the structure is fixed. However, in distributed intelligent systems one would expect the structure to be variable so that resources are used efficiently in response to the needs of a specific task. For example, in a human organization, it may be necessary to communicate information to another member only when the latter needs it. Unnecessary communication and the transfer of voluminous irrelevant data can cause severe degradation of performance by increasing the cognitive load of the individual members.

One way of introducing variability in the structure of an organization is to control the interactions between the decisionmakers through the settings of the switches in the SA and RS transitions. Consider the two person organization shown in Fig. 5., in which two patterns of interactions are allowed (Monguiller, 1987):

Setting 1 DM1 receives the situation as assessed by DM2 in his IF stage; he issues a command to DM2 and gives his own response in his RS stage. The branch (2) of switch s_1 and the branch (1) of s_2 are always chosen.

Setting 2 DM1 and DM2 have completely parallel activities in their treatment of the input, i.e., the branches (1) of s_1 and (2) of s_2 are always chosen. DM1 and DM2 never interact.




Fig. 5. A variable structure organization.

Unfortunately, the representation of variable structure organizations with ordinary Petri Nets and with switches to control the variable interactions is inadequate for two reasons:

- (a) The decision rules of the switches have to be correlated, since only two out of the four combinations of active branches are allowed for any incoming input: if (u, v) represents the numbers of the branches of the switches s_1 and s_2 , respectively, which are activated at each time, then only branches (2, 1) for Setting 1, and (1, 2) for Setting 2 are valid.
- (b) The Information Fusion stage of DM1 or the Command Interpretation stage of DM2 have different rules of enablement depending on the chosen switch setting. In the representation of Fig. 5, the transition modeling the IF stage of DM1 is enabled when there is a token in each of its two input places. But DM1 has no way of knowing when he will receive some information from DM2. If DM2 sends to DM1 either some information, or a null message to tell DM1 to continue his processing, the deadlock may be overcome: but all the advantage of variability has been lost, since either DM1 has to wait before continuing the processing or null messages must be sent.

This simple example shows that the representation of variable organizations must take into account the two requirements: *correlation of rules* and *deadlock avoidance*. A possible solution is to introduce switches in the IF and CI stages, and to associate with the Petri Net depicting the organization (where the branches of the switches have been labeled by integers), a table

showing the intercorrelation of the switch rules. However, this approach leads very quickly to a complicated representation of the organization by a Petri Net and a large table of switch settings; the transparency of the Petri Net formalism is lost.

The preferred approach is to use an extension of ordinary Petri Nets known as Higher Level Petri Nets (Genrich and Lautenbach, 1981), or, in this case, as Predicate Transition Nets. In the extension, tokens have identity, places have predicates associated with them, connectors are labeled with the tokens they can carry, and transitions have logical formulas attached to them that determine whether the transition is enabled or not. The application of Predicate Transition Nets to the modeling of variable structure organizations was done recently by Monguillet (1987).

The rules that control the interactions and, therefore, determine which structure is appropriate at any time can be of any kind. However, it is possible to distinguish three types of variability on the basis of the cause for change. The organization adapts its structure to changes in:

- (1) the input it processes.
- (2) the environment.
- (3) the system's parameters.

The performance of a system may degrade strongly when individual components are affected; for example, the loss of a communication link in an organization with a fixed structure may very well mean that deadlock will occur in the flow of information and the organization may cease to function.

These three different types of variability can be related to the properties of *Flexibility* and *Reconfigurability*. Flexibility means that the organization can adapt to the tasks it has to process, or to their relative frequency of occurrence. Reconfigurability means that it can adapt to changes in its resources or in its objectives. The flexibility and reconfigurability of real human organizations is another indication of the intelligence resident in the system. Clearly, the selection of the proper structure is not done through classical rational decisionmaking.

Given, now, that organizations with fixed structures can be generated algorithmically and that organizations with variable structures can be described precisely and rigorously, the question of evaluation or assessment arises.

MEASURES OF PERFORMANCE

The organization has been considered as a distributed intelligence system which must perform a certain task in response to a set of inputs. This set has been modeled by a finite alphabet X where the input x takes its values with a discrete probability distribution $p(x)$. The organization produces a response y , and there is a cost function J for the evaluation of the response.

In this context, one can define different measures of performance (MOP's) in order to assess the effectiveness of the organization in performing its task :

Accuracy

This measure of performance, J , evaluates how well the organizational responses correspond to the desired responses (Andreadakis and Levis, 1987). For each input x , there exists an ideal or desired response y_d . When the organization produces y , then $c(y, y_d)$ is a function that assigns a cost to the discrepancy between the desired and the actual response. When the algorithms used are deterministic, there is one response y provided for each input x ; otherwise, several responses can be given for the same input x at different times. The actual response also depends on the decision strategy used by each decisionmaker.

A pure organizational strategy is a strategy in which the DMs choose always the same algorithm in their Situation Assessment stage and, for each possible value of the input of their Response Selection stage, the same algorithm to produce a response. If there are $n(s)$ switches s_i in the entire organization, if the alphabet of inputs of the switch s_i has n_i terms, and if switch s_i has U_i branches, then the maximum number of pure organizational strategies is the following:

$$n(\text{pure}) = \prod_{i=1}^{n(s)} [U_i]^{n_i}$$

with the assumption that n_i is equal to 1 whenever the probabilities of the corresponding switch are not conditioned, as is the case for the SA stage. This number is the maximum number of possible pure strategies, because the probabilities are conditioned on the inputs z_j of the switch. The z_j are themselves outputs of other algorithms and may not describe their whole alphabet. For instance, if the value $z_j = z_j'$ is never reached, the pure organizational strategies for which:

$p(u = i \mid z_j = z_{j'}) = 1$
are never used.

The accuracy of the organization for the given strategy δ is then defined by:

$$J(\delta) = \sum_i p(x_i) \sum_j c(y_j, y_d) p(y_j \mid x_i)$$

The evaluation procedure for J is shown in Fig. 6.




Fig. 6. Model for evaluating the accuracy measure.

Timeliness

The timeliness of the response, T , measures the extent to which these responses are provided at the right instants. It is possible to define several measures of timeliness. It can be, for example, the probability that the response time lies within a certain interval $[T_{\min}, T_{\max}]$. It can be also the expected delay, which is defined as follows for each organizational strategy: for each input x_i , the time delay $r(x_i)$ for processing is computed. The expected delay is then determined by:

$$T(\delta) = \sum_i p(x_i) r(x_i)$$

These two measures are appropriate for all organizations. But in considering them as distributed intelligent systems, three more measures can be introduced (Grevet, 1988): *synchronization, information consistency, and coordination*.

Synchronization

The concept of synchronization in decision-making organizations is related to the fact that some of the activities that take place are asynchronous and concurrent. Therefore, the various decision-makers work in parallel and interact at instants that are not predetermined. The protocols of interaction can take various forms: for example, DM_1 may be requested to always wait for instructions from DM_2 before proceeding with any further action; on the other hand, DM_3 may be allowed to decide whether or not he will wait for these instructions depending on the context of tasks.

Synchronization is related to biases due to the value of the information when the decisionmaker actually processes it. (The use of old data may cause a degradation of performance; while waiting for the updated information may also cause performance degradation.) Therefore, an organization is perfectly synchronized, for the whole decision-making process, when the decision-makers do not have to wait to receive the information that they need.

A quantitative measure of synchronization is the expected value of the sum of the maximum delays at each process over all the processes in the organization for all inputs x . The delays can be computed from the analysis of the Petri Net of the organization: these delays are the extra time that it takes for a transition to be enabled because of the unavailability of a needed token from another decisionmaker.

However, perfect synchronization of the organization does not imply necessarily that the delay for the processing of one input will be low. Two decisionmakers can be perfectly synchronized but very slow; they can also be not perfectly synchronized, but very fast. In the same way, good synchronization will not ensure that the organizational response will be perfectly accurate: indeed, in order to be well synchronized, one decisionmaker might use procedures (or algorithms) which introduce important biases in his response.

Synchronization characterizes, in a way, the timeliness of the organization. It provides some

information on the evolution of the decisionmaking process and on the dynamics of the interactions.

Consistency of Information

Consistency of information assesses the extent to which different items of information can be fused together without contradiction. Therefore, it is essentially task dependent. One can evaluate the extent to which two data are inconsistent, if the subsequent processing or actions that will take place for each item of information are not the same. Thus, one cannot speak of the inconsistency of information in general, but must refer to the context and the task at hand.

If the distributed sensors provide different data about the same events, or if different organization members arrive at different situation assessments, the inconsistency must be overcome for the organization to provide a coherent response. Decision aids, such as expert systems, can be used to resolve these conflicts (Perdu, 1987).

A quantitative measure of information consistency has been developed, again based on the description of the organization by Predicate Transition Nets. Since in these nets tokens have identity, they can be marked by their arrival time, the task to which they belong, and some other attribute that characterizes the data they contain - they can belong to different classes. The degree of information consistency for a transition measures the extent to which the tokens enabling that transition are of the same class; the expected value of this measure over all transitions and all inputs is the organization's measure of information consistency. Clearly, this measure is related to accuracy; the latter, however, considers only the final output.

Coordination

Coordination has been defined as a composite measure. The firing of a transition (or the execution of a process) is coordinated, if it is consistent and synchronized. Consequently, the execution of a task is coordinated if its component processes are coordinated, and an organization is coordinated, if and only if it is coordinated for all the tasks performed.

All the values of these measures are obtained for a specific organizational strategy. As the whole

space of decision strategies is spanned, the corresponding values of the measures are obtained. This can be interpreted as a mapping from the strategy space to the performance space. The locus of points in the performance space - the generalized performance workload space - characterizes the particular organizational design. By comparing qualitatively and quantitatively the loci of different organizations with respect to each other and against requirements, evaluations and trade-off analyses can be carried out. The important aspect is that the locus represents the mapping of all admissible strategies since the rules by which intelligent decisionmakers may use to select strategies are not known. The procedure can be refined by allocating different probabilities to the use of different strategies to reflect decisionmaking styles and preferences. However, this has limited usefulness until relevant data are obtained.

RESULTS

The value in the development of a mathematical theory of organizations is twofold: to analyze organizations and discover phenomena and behaviors that can be tested in experiments and that can provide insight in the causal relationships that govern them, and to design organizations that can meet, with predictability, performance requirements.

At this stage in the development of the theory, results have been obtained that have either provided plausible explanations for behavior observed in practice, or are appropriate for the formulation of experimental hypotheses. Some of these results will be described, but only qualitatively. The complete analyses are available in the literature.

Fixed and Variable Structure Organizations

First results on the comparison of fixed and variable structure organizations have confirmed what has been expected. A variable structure organization can perform well over a wide variety of conditions. But, for specific conditions, one can design a fixed structure organization that can perform better. Or, as seen from the requirements point of view, one can design a fixed structure decisionmaking organization (FDMO) that can meet accuracy (J) requirements, or timeliness requirements (T), but one may not be able to design one that meets both types of requirements. However, it is possible to design a variable structure one that meets both requirements. But the flexibility inherent in the VDMO is obtained at a cost. The VDMO may not be able to achieve the best accuracy achievable by an optimized FDMO, or the best speed of response possible by

another FDMO. This is shown clearly in Fig. 7, where the performance loci of three organizational designs are shown (Monguillet, 1987). FDMO1 optimizes speed of response at a sacrifice in accuracy; FDMO2 optimizes accuracy but at a cost in time. The VDMO achieves intermediate levels of accuracy and speed of response. It can never achieve the highest accuracy possible or the fastest response possible by FDMO2 and FDMO1, respectively, but does not attain the worst values either. Unfortunately, the flexibility inherent in variable structure organizations is achieved through an increase in cognitive workload. A higher-level cognitive task has been introduced - the management of the variability.




Fig. 7. Comparison of fixed and variable structure organizations.

This phenomenon that has appeared again and again in the analysis of distributed decisionmaking organizations has two consequences. One is psychological and one is organizational. The first one can be described by the term *metadecision making*. The second is the introduction of intelligent decision aids.

Metadecisions

In addition to the decisions that pertain directly to the task for which the human decisionmakers are presumably well trained, they also have to make decisions about how to decide, or, as Einhorn and Hogarth (1981) have described it, deciding how to choose. This metadecision making, in addition to introducing a cognitive processing "overhead," has often had a deleterious effect on performance, not only because scarce human cognitive resources are used, but also because it shifts the domain of decisions from the one in which the DM is an expert one to one in which he may be less well trained or less expert (Weingaertner and Levis, 1988)

Expert Systems

One approach that is often proposed to address this problem is the introduction of decision support systems, especially those based on the expert system model in artificial intelligence. The rationale is obvious: human and artificial (or machine) intelligence are now distributed within the organization with the expected result that higher performance can be achieved or more stringent requirements met.

In some recent work, Perdu (1987) modeled expert system decision aids using Predicate Transition Nets. These models were integrated in the Petri Net description of the organization and the overall design could then be evaluated in a consistent manner. The case of two decisionmakers carrying out a data fusion task was considered. Three basic strategies were investigated. DM1 ignores the information provided by DM2; DM1 uses a weighted estimate of the two assessments by taking into consideration the way data used to make these assessments were obtained by each DM; and DM1 uses an expert system. Two implementations of the expert system were modeled; one based on boolean logic and one on fuzzy logic. The results are shown in Fig. 8. Again, it is clear that improved accuracy can be obtained at the cost of longer response time. Therefore, if the response time requirements are too stringent, the expert system will be useless because too much time is needed to perform the data fusion task.




Fig. 8. Accuracy and timeliness trade-off.

Migration of Control

Another aspect of distributed intelligent system behavior that appears in organizations is the apparent migration of control (Kahne, 1983). In distributed decisionmaking organizations, one can study the effect of each decision maker on the organization's performance. This information is readily obtained by considering isoquants on the performance locus: for example, if the decision strategies of all decisionmakers but one are held constant, then the locus of performance when the strategies of the last one only are varied represents the effect that DM has on performance. (This is the way the software that has been developed constructs the performance loci). By observing the curvature of the loci, one can determine both qualitatively and quantitatively the decisionmakers whose strategy choice affects most the organization's performance.

Changes in the organization's structure through the addition of resources - access to data bases or decision support systems - cause a change in the locus and, consequently, on the sensitivity of the performance measures to different decisionmakers' actions.

Dynamic Phenomena

The effects that have been described thus far characterize the steady state performance of distributed decisionmaking organizations. Some very interesting dynamic phenomena have also been observed and analyzed recently with the use of a simulation system based on Predicate Transition Nets (Grevet et al., 1988). It has been observed in simulations that the introduction of additional resources to aid a specific decisionmaker may cause the synchronization between DMs to be degraded with an attendant reduction in performance. Or, when different DMs use different protocols, individual tasks may become blocked at a specific decisionmaker and never be processed, while other later tasks are performed. This is not a system deadlock; the phenomenon appears only when the tokens are given identities that affect the execution of the protocols.

It is also becoming clear from the analyses that changes in the protocols, while preserving the basic pattern of interactions and the processing resources, can yield significant improvements of performance. Conversely, changes in the structure without well analyzed changes in the protocols and the resources may cause performance degradation. Both effects are well known in the operational world of organizational design. Thus, the ability to design protocols to meet performance requirements on the organization is seen as a key area where theoretical work needs to be done.

CONCLUSIONS

Human organizations are indeed distributed intelligent systems. The study of these organizations and the development of mathematical theories for their analysis and design not only is of interest to organization designers, but also to the designers of physical distributed intelligent systems. The theory, even in its early stages of development, is showing that distributed intelligent systems can exhibit a wide variety of not well understood behaviors, both in steady state and in transient state. Furthermore, the understanding of the properties of systems in which both human and machine intelligence co-exist and interact should be of high priority, since these are the real systems that are being designed and implemented.

REFERENCES

- Andreadakis, S. K., and A. H. Levis (1987). Design methodology for command and control organizations. *Proc. 1987 Symposium on C² Research*, National Defense University, Washington DC.
- Boettcher, K. L., and A. H. Levis (1982). Modeling the interacting decisionmaker with bounded rationality. *IEEE Trans. on Sys. Man and Cybernetics, SMC-12*, 334-344.
- Drenick, R. F. (1986). *A Mathematical Theory of Organizations*. North-Holland, New York.
- Einhorn, H. J., and R. M. Hogarth (1981). Behavioral decision theory: Processes of judgement and choice. *Annual review of Psychology*, 32.
- Galbraith, J. R. (1977). *Organization Design*. Addison-Wesley, Reading, MA.
- Genrich, H. J., and K. Lautenbach (1981). System modeling with high-level Petri Nets. *Theoretical Computer Science*, 13, 109-136.
- Grevet, J. L. (1987). Decision aiding and coordination in decisionmaking organizations. LIDS-TH-1737, Laboratory for Information and Decision Systems, MIT, Cambridge MA.
- Grevet, J. L., L. Jandura, J. Brode, and A. H. Levis (1988). Execution strategies for Petri Net simulations. *Proc. 12th IMACS World Congress on Scientific Computation*, Paris.
- Kahne, S. (1983). Control migration: a characteristic of C³ systems. *IEEE Control Systems Magazine*, 3, 15-19.
- Keene, P. G. W., and M. S. Scott Morton (1978). *Decision support systems: An organizational perspective*. Addison-Wesley, Reading, MA.
- Levis, A. H. (1984). Information processing and decision-making organizations: A mathematical description. *Large Scale Systems*, 7, 151-163.
- Louvet, A. C., J. T. Casey and A. H. Levis (1987). Experimental investigation of the bounded rationality constraint. *Proc. 1987 Symposium on C² Research*, National Defense University, Washington DC.
- Minsky, M. (1986). *The society of the mind*. Simon and Schuster, New York.
- Monguillet, J.-M. (1987). Modeling and evaluation of variable structure organizations. LIDS-TH-1730, Laboratory for Information and Decision Systems, MIT, Cambridge MA.
- Peterson, J. L. (1981). *Petri Net theory and the modeling of systems*. Prentice-Hall, Englewood Cliffs, NJ.
- Perdu, D. M. (1987). Modeling and evaluation of expert systems in decision making organizations. LIDS-TH-1726, Laboratory for Information and Decision Systems, MIT, Cambridge MA.

Remy, P. A., and A. H. Levis (1988). Algorithmic generation of feasible system architectures. *Proc. 12th IMACS World Congress on Scientific Computation*, Paris, June.

Remy, P. A., A. H. Levis, and V. Y. Jin (1988). On the design of distributed organizational structures. *Automatica*, 24, 81-86.

Weingaertner, S. T., and A. H. Levis (1988). Evaluation of decision aiding in submarine emergency decisionmaking. *Proc. 3rd IFAC/IFIP/IEA/IFORS Conf. on Man-Machine Systems*, Oulu, Finland.

END
DATE
FILMED

6-1988

DTIC